Aerospace Science and Technology 135 (2023) 108185

Contents lists available at ScienceDirect

Aerospace Science and Technology

journal homepage: www.elsevier.com/locate/aescte

Multi-fidelity aerodynamic design and analysis of propellers for a heavy-lift eVTOL

Tao Zhang [a,1], George N. Barakos [a,*], Furqan [b], Malcolm Foster [c]

[a] School of Engineering, University of Glasgow, Glasgow G128QQ, Scotland, UK
[b] Future Flights, GKN Aerospace, GKN Global Technology Centre, Taurus Rd, Patchway, Filton, Bristol BS34 6FB, England, UK
[c] Special Projects, GKN Aerospace, GKN Global Technology Centre, Taurus Rd, Patchway, Filton, Bristol BS34 6FB, England, UK

ARTICLE INFO

Article history:
Received 18 August 2022
Received in revised form 7 February 2023
Accepted 7 February 2023
Available online 13 February 2023
Communicated by Grigorios Dimitriadis

ABSTRACT

This paper presents a multi-fidelity propeller design and analysis process, and demonstrates it for the preliminary design of a heavy-lift eVTOL vehicle proposed by GKN aerospace, known as Skybus. The multi-fidelity framework integrates tools and results of variable fidelity levels. A global scan of the design space was first carried out at the low-fidelity stage. The low-fidelity output was converted to 3D shapes and grids through an automatic meshing tool, and high-fidelity CFD simulations and gradient-based optimisation were launched to deliver further improved designs. The method was first verified using a benchmark multi-modal test function. The framework was then demonstrated for the propeller design of the Skybus vehicle. Rectangular blade designs with linear twist were first derived through the Blade Element Momentum Theory and used as the input to the low-fidelity stage for simplicity. Upon the baseline design, the low-fidelity stage managed to find an improved shape with 3% reduced power, while the high-fidelity optimisation further reduced the power by 2.2% under the equality thrust constraint. In addition, this work also reports the propeller pitch-RPM performance map tool that has effectively supported the Skybus development through fast performance predictions and operating condition determination. For propellers with both pitch and RPM regulation, the performance map explicitly correlates various performance scopes and to aid optimisation. The multi-fidelity construction of the performance map and demonstrations of its usage are then presented.

1. Introduction

Electrical Vertical Take-off and Landing (eVTOL) vehicles have received strong attention from industry and academia, and have been widely conceived as the ideal choice for Urban Air Mobility (UAM) [1]. To assess the feasibility of rarely explored heavy-lift eVTOL designs, GKN aerospace recently proposed the Skybus concept with a maximum capacity of 30 passengers, for short-distance transport in densely populated cities. To bridge high-speed and low-speed capabilities, the Skybus vehicle was designed as a compound rotorcraft [2,3] integrating multiple lifting surfaces and six propellers with convertible axes. The propellers are to be powered by electrical motors, and will have both pitch and RPM controls. In low-speed flight, e.g. hover or taking-off/landing, the entire vehicle weight is carried by the propellers. This requirement poses challenges for the propeller design. More advanced and efficient tools are necessary to accelerate the design and analysis process.

At the preliminary propeller design stage, tools and methods of various fidelity levels are widely used. Typically, low-fidelity methods are associated with high uncertainties and low computational costs, while high-fidelity methods are associated with low uncertainties and high computational cost. The concept of multi-fidelity modelling was hence proposed to bridge the benefits of low- and high-fidelity methods, and allow their drawbacks to be mutually compensated. Modern multi-fidelity modelling have been focusing on models joining low- and high-fidelity results, and the difficulty in many methods is to strike to balance among accuracy, cost, and complexity. Regression-based multi-fidelity surrogate models, such as co-Kriging [4] and Multi-fidelity Gaussian Regression [5], have been popular due to the simplicity and high technology readiness. There have been a few attempts in applying such methods to the design and analysis of marine propellers [6], high-altitude propellers [7], and ducted propellers [8]. In these studies, multi-fidelity models were incorporated into the design optimisation to

* Corresponding author.
E-mail addresses: t.zhang.4@research.gla.ac.uk, tz77@leicester.ac.uk (T. Zhang), george.barakos@glasgow.ac.uk (G.N. Barakos), furqan@gknaerospace.com (Furqan), Malcolm.Foster@gknaerospace.com (M. Foster).
[1] Currently at the University of Leicester.

https://doi.org/10.1016/j.ast.2023.108185

Nomenclature

Latin

c	Blade Local Chord ..... m
I	Objective Function
N	Number of Samples
P	Power ..... W
p_{ref}	Acoustic Reference Pressure, $p_{ref} = 2 \times 10^{-5}$ [Pa]
Q	Torque ..... N m
r	Blade Local Span Position ..... m
R	Blade Radius ..... m
R	Residual Vector
RPM	Revolutions Per Minute
SPL	Sound Pressure Level, $OSPL = 20log_{10}(\frac{p'}{p_{ref}})$
T	Thrust ..... N
W	Conservative Flow Variables
X	Sample or Predication Set
$\hat{X}$	Data Set Mean Value, $\hat{X} = \frac{\Sigma(X_i)}{N}$

Greek

α	Design Variables
β	Blade Local Pitch Angle ..... °
λ	Adjoint Vector
$\sigma(X)$	Standard Deviation, $\sigma(X) = \sqrt{\frac{\Sigma(X_i - \hat{X})^2}{N}}$
$\delta(X)$	Prediction Relative Error, $\delta(X) = \frac{X_{prediction} - X_{sample}}{X_{sample}} \times 100\%$

Subscripts and superscripts

x, y, z	3D Cartesian Directions
i, j, k	3D Block Index

Acronyms

CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics
eVTOL	Electric Vertical Take-off and Landing
HMB3	Helicopter Multi-Block 3
IDW	Inverse Distance Weighting
SST	Shear Stress Transport
RANS	Reynolds Averaged Navier Stokes
UAM	Urban Air Mobility

provide fast and accurate predictions. However, as the problem dimensionality increases, the required number of sampling points to construct predictive models grows dramatically. Even with the help of multi-fidelity modelling, the required computational cost would still be prohibitive.

In many cases, however, integrated predictive models may not be necessary to accomplish the design and analysis mission. Low-fidelity results can be simply used to point out where high-fidelity analyses should focus. This idea was explored by Molina-Cristóbal et al. for the optimisation of submarine propulsion systems [9]. A global search was applied to the low-fidelity stage, to identify potential optimal designs, while a local search was coupled with high-fidelity methods to focus on these designs to further locate the final optima. This method effectively delivered improved designs and reduced the computational cost by almost half, compared to high-fidelity methods. However, due to gradient-free nature of the employed optimisation algorithms, hundreds or thousands of high-fidelity evaluations were required.

For high-fidelity aerodynamic shape design and optimisation problems, it is well established that gradient-based methods are more efficient than surrogate-assisted, gradient-free methods [10]. Coupled with high-fidelity methods, gradient-based approaches require considerably reduced objective function evaluations, especially when dealing large amounts of design variables [10]. The gradient computation can be accomplished through adjoint formulations [11–13]. The pitfalls, however, are that the computational costs may be high and the final result may be trapped in local optima for multi-modal functions due to the gradient-based searching algorithm. A simple and effective solution is to link the high-fidelity method with a low-fidelity global search of the design space. The high-fidelity computational cost can be effectively cut with an improved input, and the implementation is simple enough to avoid extensive modifications of complex high-fidelity tools. This idea of combining low-fidelity modelling and high-fidelity gradient-based design and analysis is rarely explored.

This paper presents a mixed-fidelity methodology for efficient eVTOL propeller design and analysis and at the same time demonstrates the method for a demanding heavy-lift eVTOL design. The propeller design space was first scanned using low-fidelity tools. High-fidelity analyses and gradient-based optimisation were then carried upon the low-fidelity output. The method was first evaluated and demonstrated using a benchmark multi-modal test function. Demonstrations of propeller designs from scratch for the Skybus vehicle were then presented in detail. At the low-fidelity stage, 3000 potential propeller designs were sampled and quickly evaluated using the low-fidelity XRotor code [14,15]. Their performance was analysed and optimal designs were selected. The selected designs were then converted into CAD (Computer-Aided Design) models and meshed with chimera grids and an automatic meshing framework [16,17]. High-fidelity CFD (Computational Fluid Dynamics) simulations and adjoint-based optimisation were then performed using the in-house HMB3 solver [18–21]. Compared to the simple initial design, the low-fidelity stage yielded a 3% performance improvement, and the final high-fidelity-optimised design showed a further 2.2% improvement. Detailed shape comparisons between the low- and high-fidelity optimal designs are also presented. Moreover, this work further reports the propeller pitch/RPM performance map tool that successfully facilitated the Skybus development.

2. Numerical methods and tools

2.1. Multi-fidelity propeller analysis and design framework

The flowchart of the proposed multi-fidelity analysis and design framework are shown in Fig. 1. The idea is to first scan the global design space using low-fidelity tools to identify possible optimum. Once the potential optima are found, gradient-based high-fidelity methods are launched to deliver both accurate analysis and precise localisation of the optimal design in the design space. The low-fidelity and high-fidelity data was later integrated in multi-fidelity response surfaces for propeller performance analysis. This simple and straightforward multi-fidelity approach was first verified and compared with other single- or multi-fidelity approaches using a benchmark test function in later sections, and was then applied to eVTOL propeller design and analysis. In the present work, the proposed framework consists mainly of the low-fidelity and high-fidelity stages. These two stages are bridged by an automatic CAD and mesh generation framework that converts the low-fidelity outputs into inputs to the high-fidelity stage.

The low-fidelity stage provides a quick, parametric scan of the design space. At this stage, in total 3000 potential designs

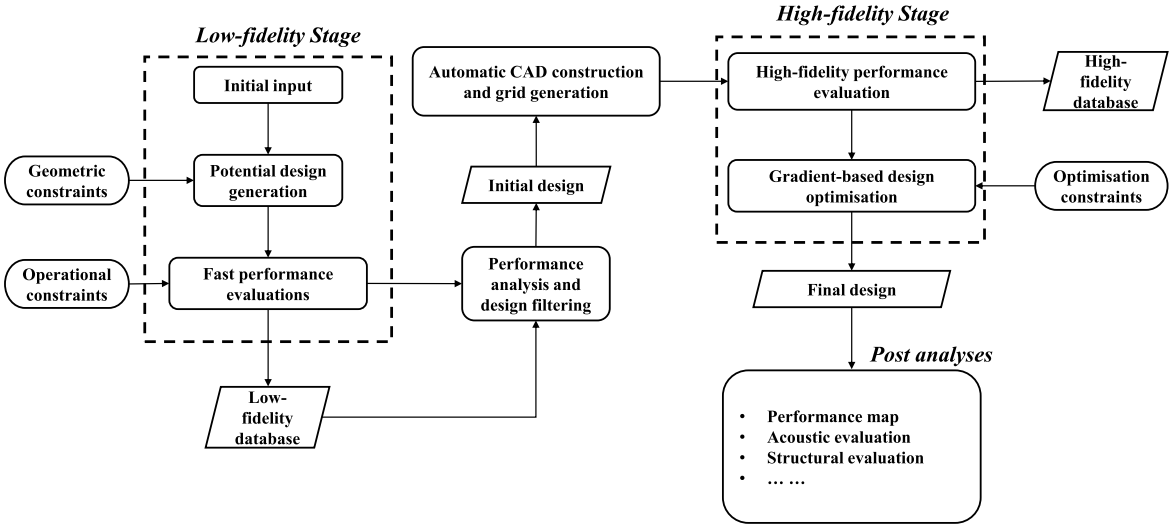

Fig. 1. Flowchart of the proposed multi-fidelity analysis and design framework for propellers.

are sampled and quickly assessed using low-fidelity methods for the case studied. The low-fidelity analyses were carried using the XRotor code [14]. The low-fidelity stage outputs an initial performance database, as well as, a low-fidelity optimal design, subject to constraints. This initial design corresponds to the global optimum within the design space resolved by the low-fidelity tool. However, it should be stressed that this optimal result is associated with high levels of uncertainty due to the low-fidelity nature of the method. Nonetheless, it can serve as a fine input to the high-fidelity stage for further analysis and design optimisation.

High-fidelity evaluations are then carried out focusing on a small region indicated by the low-fidelity modelling. The low-fidelity design output is channelled into the high-fidelity stage through an automatic geometry and mesh composition framework of the HMB3 toolkit [16,17]. High-fidelity CFD simulations were then carried out using the HMB3 solver to evaluate the performance of the low-fidelity design input. High-fidelity gradient-based shape optimisation is launched to further improve the propeller design subject to constraints. Typically, gradient-based optimisation only offers local optima close to the initial input. In this case, however, a global scan of the design space was first conducted at the low-fidelity stage. Although the results were associated with large uncertainty, they can help avoid obvious local optima in the design space and lead to improved solutions.

The low-fidelity results narrow the design space to be explored by the high-fidelity methods, reducing the computational cost. Note that surrogate models and optimisation algorithms are considered unnecessary for the low-fidelity stage given the high uncertainty of the method, at this particular stage. The low-fidelity optimum can be obtained by simply sorting the samples. The high-fidelity analysis and optimisation then eliminates the uncertainty associated with the low-fidelity methods.

The proposed framework also features high-levels of automation and versatility. The framework handles the sampling, computation, data processing/conversion, and optimisation mostly automatically. Only minor user interventions and decision making are necessary. The analysis tools and methods, including the parametrisation and optimisation, can be easily replaced, if necessary. In addition, databases are generated by this framework, and can be recycled for further analysis through multi-fidelity response surfaces.

2.2. HMB3 flow solver

As a high-fidelity flow solver and the core of the high-fidelity analysis stage, the in-house Helicopter Multi-Block 3 (HMB3) [12,22] CFD code was used in the present work. The code has been widely used in simulations of rotorcraft flows [23–26]. HMB3 solves the Unsteady Reynolds Averaged Navier-Stokes (URANS) equations in integral form using the Arbitrary Lagrangian Eulerian (ALE) formulation for time-dependent domains, which may include moving boundaries. The Navier-Stokes equations are discretized using a cell-centred finite volume approach on a multi-block, structured grid:

$$\frac{d}{dt}\left(\mathbf{W}_{i,j,k}V_{i,j,k}\right) = -\mathbf{R}_{i,j,k}\left(\mathbf{W}_{i,j,k}\right), \tag{1}$$

where i, j, k represent the cell index, $\mathbf{W}$ and $\mathbf{R}$ are the vector of conservative flow variables and residual respectively, and $V_{i,j,k}$ is the volume of the cell i, j, k. In the present work, simulations of propellers in hover and axial flight were performed in the non-intertia Rotating Reference Frame (RRF) method [19,20,27] in HMB3. This converted the modelling to steady simulations with rotational periodicity and enabled the use of adjoint computations for gradients. The high-fidelity simulations were performed on grids with 6.3 million cells for a single periodic domain, which was more than enough to resolve the dominant flow features according to previous studies [20,27]. A slightly coarser mesh of 4 million cells were used for the adjoint-based optimisation to reduce the computational cost, but the optimisation results were later verified using the 6.3-million-cell grids.

In the present work, near-field propeller acoustic data was extracted directly from the CFD solution. The conversion of solutions between rotating and fixed frames allowed unsteady and periodic pressure fluctuations to be extracted at microphone positions [20,28]. The 6.3-million-cell grids used ensured at least 10 mesh points for the wavelength up to the fourth blade passing frequency, within 3 blade radii from the centre of rotation.

To evaluate the convective fluxes, Osher-approximate Rieman solver is used, while the viscous terms are discretized using a second order central difference scheme. The 3^{rd} order MUSCL (Monotone Upstream-centred Schemes for Conservation Laws) approach was used to provide high-order accuracy in space. The chimera/overset grid method [29] was used in this work. In the

present work, simulations were performed with the $k-\omega$ SST (Shear Stress Transport) [30] turbulence model.

The HMB3 solver is also equipped with an automatic grid generation framework [16,17]. The automation framework is capable of delivering composed geometries and high-quality, ready-to-run structured grids for various purposes, upon simple definitions or existing CAD models, with minor human interventions. It greatly eases the efforts required for the preparation of high-quality grids for numerical analysis. This framework was used in the present work to bridge the low-fidelity and high-fidelity stages.

2.3. HMB3 adjoint solver and gradient-based optimisation

The adjoint method [11,12] is known for its efficient handling of hundreds of thousands of design variables with a handful of objective functions. This is especially suitable for the current optimisation study of propeller blades, which involves dozens of design variables governing blade shapes with just a few objective/constraint functions e.g. the overall thrust and torque.

In the context of CFD analysis, a typical objective function I can be explicitly defined as $I(\mathbf{W}(\alpha),\alpha)$, which is a function subject to the design variable vector α and the flow variables $\mathbf{W}(\alpha)$. The flow variables still satisfy the governing equation in the form $\mathbf{R}(\mathbf{W}(\alpha),\alpha)=\mathbf{0}$.

The core of many gradient-based optimisation problems, as also involved in this work, is to solve the sensitivity equation of I relative to the design variables α as follows:

$$\frac{dI}{d\alpha}=\frac{\partial I}{\partial \alpha}+\frac{\partial I}{\partial \mathbf{W}}\frac{\partial \mathbf{W}}{\partial \alpha}. \tag{2}$$

The difficulties solving this equation lie in the second term on the RHS, especially the term $\frac{\partial \mathbf{W}}{\partial \alpha}$. It can be implicitly derived from the governing flow equations, but the computational requirement would be excessive when using direct differentiation or finite differences because of the large amount of design variables.

Alternatively, introducing an adjoint vector λ correlating the cost function and the governing equations, the sensitivity equation (2) can be recast in the adjoint form as:

$$(\frac{\partial \mathbf{R}}{\partial \mathbf{W}})^T\lambda=-(\frac{\partial I}{\partial \mathbf{W}})^T. \tag{3}$$

$$\frac{dI}{d\alpha}=\frac{\partial I}{\partial \alpha}+\lambda^T\frac{\partial \mathbf{R}}{\partial \alpha}. \tag{4}$$

The major effort solving the sensitivity equation aligns with solving the adjoint vector λ coupled in the linear system in Equation (3). This linear system scales with the number of objective functions and is irrelevant to the design variables. This property makes the adjoint formulation especially suitable for aerodynamic shape optimisation. Details of the HMB3 adjoint formulations can be found in previous optimisation studies of rotors, wings, and ducted propellers [12,19–21].

With the gradients computed, design changes can be determined to improve the objective performance subject to constraints. The flow chart of the gradient-based optimisation framework coupled with HMB3 is presented in Fig. 2. The Sequential Least-Square Quadratic Programming (SLSQP) algorithm [31] as provided in the *NLopt* library [32] is used as the optimiser in this study. The SLSQP algorithm approximates and solves the objective functions with a sequence of least-square/quadratic programming problems. The constraints are approximated with linear functions. This algorithm is frequently used due to its effectiveness and efficiency. Based on the initial objective, constraint, and gradient information, the optimiser solves for the searching directions in the design space and step sizes. The design variables are then modified accordingly, and new CFD computations are launched to acquire the new cost function value and gradients. The iterations continue until the gradients or step size are approaching zero, or the maximum iteration has been reached. An advanced mesh deformation algorithm is implemented in HMB3 based on Inverse Distance Weighting (IDW) interpolations [19,33,34] to bridge the shape deformation and the adjoint system. More details of the optimisation framework can be found in Refs [13,20].

2.4. XRotor code

The XRotor code is an open-source tool developed by Drela et al. [14,15,35] for performance predicting and design of propellers. The code is based on extended classic blade-element/vortex theories, in combination with lifting line and panel methods. It is capable of quickly predicting, or matching propeller performance for specific geometries, as well as solving inverse design problems. Nonetheless, the code relies heavily on accurate blade sectional aerodynamic inputs. Due to inherent limitations of the methods, it is also incapable of handling complex flow phenomena e.g. stall, flow separation, and strong compressibility.

In the present work, the blades were modelled by lifting lines coupled with a prescribed vortex formulation. The computation applied empirical tip-loss, viscous, and compressibility correction factors, but did not account for detailed boundary layer and 3D effects. The required aerofoil characteristics were computed using the XFOIL code [36]. For simplicity, the characteristics were considered only functions of the angle of attack. The XRotor is used in this work for as the low-fidelity tool for quick evaluation of the aerodynamic performance of 3000 propeller designs, but the delivered results are to be corrected by high-fidelity methods, later in the analyses.

2.5. Kriging response surfaces

The Kriging surrogate model [37] is a spatial interpolation method based on Gaussian regression. Note that the proposed multi-fidelity propeller design framework doest not need surrogate models as the optimisation is primarily gradient-based. The current work adopted the Kriging model as a tool for propeller performance visualisation and interpolation/extrapolation. The Kriging model predicts the value of an unobserved evaluation point using a predictor function plus a small, stochastic variance as follows

$$Z(\mathbf{x})=Z_0(\mathbf{x})+\epsilon(\mathbf{x}), \tag{5}$$

where $Z(\mathbf{x})$ is the prediction at the unknown location $\mathbf{x}$, and $Z_0(\mathbf{x})$ is a regression of the data set. $\epsilon(\mathbf{x})$ is a realisation of a stochastic process with zero mean defined by prescribed spatial correlation among the dataset. The current work adopts the Ordinary Kriging with regression Z_0 denoted by a constant (the mean value) and $\epsilon(\mathbf{x})$ assuming a Gaussian distribution. The determination of the stochastic parameters used the maximum likely hood estimation. A nugget value of 2.2×10^{-14}, which is just 100 times of the machine epsilon of a 64-bit float number, was used to improve the numerical stability.

A multi-fidelity Kriging model [38] was also adopted in the present work. The model uses an extra Kriging process to integrate data from different fidelity levels as

$$Z_{high}(\mathbf{x})=\rho_0(\mathbf{x})Z_{low}(\mathbf{x})+\epsilon_1(\mathbf{x}), \tag{6}$$

where $Z_{high}(\mathbf{x})$ is the final, high-fidelity prediction, Z_{low} is the Kriging prediction based on the low-fidelity data, $\rho_0(\mathbf{x})$ is a correlation factor, and $\epsilon_1(\mathbf{x})$ is an addition correction. For consistency, the present work adopted the same functional and parameters for

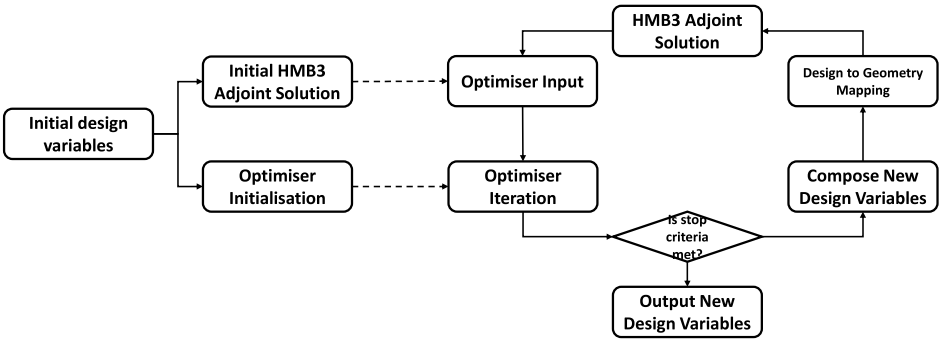

Fig. 2. Flow chart of the gradient-based optimisation framework [20].

Table 1
Test matrix of the assessment of different numerical design and optimisation methods. Method M3 corresponds to the proposed multi-fidelity approach.

Test Methods	Modelling fidelity level	Optimisation method
M1	high	high-fidelity gradient-based
M2	low+high	multi-fidelity surrogate + Bayesian
M3(present)	low+high	low-fidelity surrogate + high-fidelity gradient-based

Kriging and multi-fidelity Kriging models, e.g. Gaussian assumptions for $\epsilon_1(\mathbf{x})$ and 2.2×10^{-14} as the nugget value. The correlation factor ρ_0 assumed a constant function. These were all implemented through the open-source Surrogate Modelling Toolbox (SMT) [39].

3. Methodology demonstration and verification

This section presents a brief demonstration of the proposed multi-fidelity approach to verify its benefits over conventional methods. The benchmark multi-modal Styblinski-Tang function [40] in 2D, as illustrated in Fig. 3, was used as the objective function. The function definition is as follows:

$$f(x_i) = \frac{1}{2}\Sigma_{i=1}^{d}(x_i^4 - 16x_i^2 + 5x_i), \tag{7}$$

where $x_i \in [-5, 5]$. d represents the dimensionality and in this work $d = 2$. As can be noted from Fig. 3, this function has several local optima where the gradients are zero. However, there is only one global minimum $f_{min} = -39.166d$ at $x = (-2.9035, \ldots, -2.9035)$. For this test, the exact function was used as the high-fidelity model, as shown by the contour in Fig. 3. A low-fidelity model of the exact function was constructed using a second-order Inverse Distance Weighting approximation, as illustrated by the iso-lines in Fig. 3 (dashed lines denote negative values, and '+' denotes the sampling positions).

The methodology test matrix is presented in Table 1. The present work compares classic gradient-based approach (M1), multi-fidelity surrogate-based approach (M2), and the proposed multi-fidelity gradient-base approach (M3). All methods were implemented to find the minima of the 2D Styblinski-Tang function as shown in Fig. 3. The SQLSP algorithm [31] was employed for gradient-based searching, and the Efficient Global Optimisation (EGO) [41] algorithm was used for gradient-free, Bayesian-type optimisation. The Kriging model was used for uni-fidelity surrogate approximations. For the multi-fidelity data in M2 in Table 1, the multi-fidelity Kriging model [38] was adopted. The present study aims to evaluate the convergence, cost, and final optimal solution.

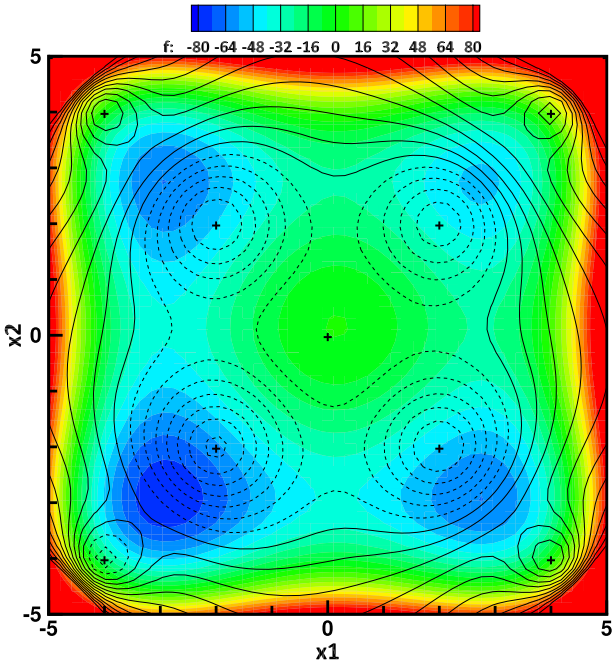

Fig. 3. Two-dimensional, multi-modal Styblinski-Tang [40] test function. Contours represent the high-fidelity model (exact function). Iso-lines represent the low-fidelity model. (Inverse Distance Weighting approximation using the '+' points. Dashed lines denote negative function values.) (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

Fig. 4(a) presents the convergence history in terms of the objective value evolution for all gradient-based approaches (M1 and M2 in Table 1), while Fig. 4(b) shows the gradient-free convergence for M3 in terms of the Expected Improvement (EI) metric [41].

For the high-fidelity, gradient-based approach M1 in Table 1, we examined four different initial points, i.e. $(\pm 1, \pm 1)$. The coordinates in the legends of Fig. 4(a) correspond to the specific initial point. In Fig. 4(a), it is clearly shown that different initial points led to different optimal solutions, typically the nearest optima. Only the initial point $(-1, -1)$ resulted in the global optimum. With different initial points, the convergence behaviour remained similar, although required iterations were different.

Fig. 4(a) also presents the convergence history of M3 in Table 1, which corresponds to the proposed multi-fidelity approach combined with gradient-based methods. Unlike typical gradient-based approaches, M3 needed no manual selection of the initial point, and it was determined by the low-fidelity optimum. Fig. 4(a) clearly shows that M3 converged faster than the typical gradient-based M2 with any initial points tested, and the final optimal solution corresponded to the global optimum.

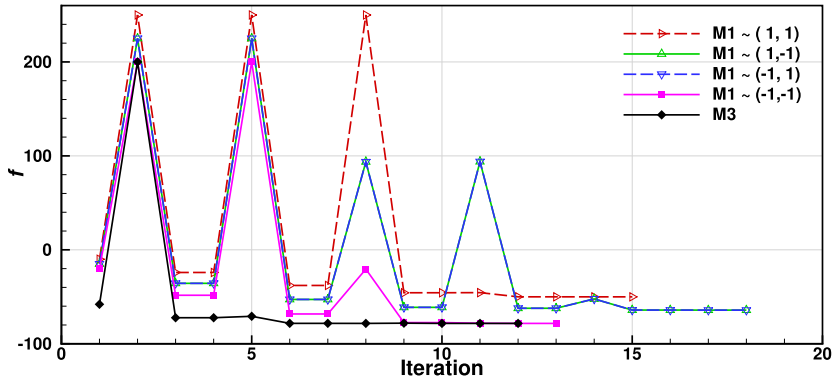

(a) Methods M1 and M3. The coordinates correspond to different initial points.

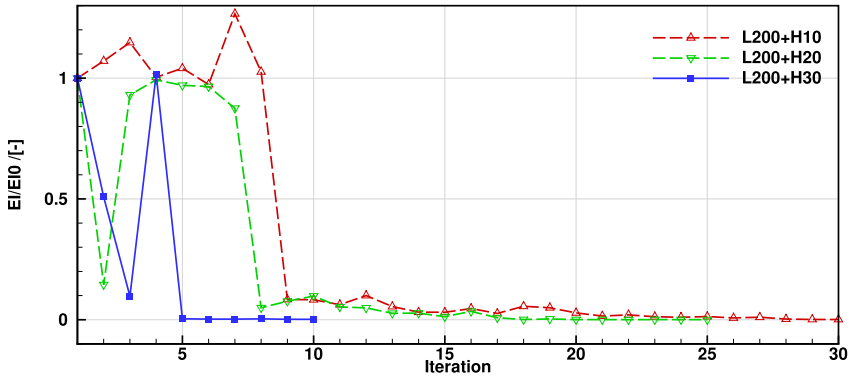

(b) Method M2. The label "La+Hb" denote that a low-fidelity samples plus b high-fidelity samples were used to generate the initial surrogate model.

Fig. 4. Convergence history of methods M1, M2, and M3 in Table 1.

Table 2
Detailed comparisons of costs and final solutions of all methods in Table 2.

Method	Initial point	Function evaluations			Total cost/[-]	Optimal solution
		Low-fidelity	High-fidelity	Gradient		
M1	(1,1)	-	12	12	24	local
	(1,-1)	-	15	15	30	local
	(-1,1)	-	15	15	30	local
	(-1,-1)	-	9	9	18	global
M2	L200+H10	200	10+25	-	37	global
	L200+H20	200	20+17	-	39	global
	L200+H30	200	30+5	-	37	global
M3 (present)	-	200	6	6	14	global

Fig. 4(b) presents the convergence history of conventional multi-fidelity surrogate-based approach M2 (Table 1). This study examined the impacts of the accuracy of the initial multi-fidelity response surface. This was realised by increasing the number of high-fidelity samples in the initial surrogate construction. The label "La+Hb" in Fig. 4(b) denotes that a low-fidelity samples and b high-fidelity samples were used to construct the initial surrogate. In the present work, a was fixed at 200, while b varied from 10 to 30. Note all samples were orthogonally generated to avoid randomness. Fig. 4(b) shows that the lower is the initial surrogate accuracy, the more iterations are needed for the convergence. Nonetheless, all cases managed to converge to the global optimum.

Table 2 presents further comparisons of all methods in Table 1 in terms of costs and final optimal solutions. The cost was estimated by assigning each high-fidelity function evaluation a dimensionless cost of 1. Each low-fidelity evaluation was assumed to have a dimensionless cost of 0.01 (in practice it may be even lower). The adjoint-based gradient evaluation had the same cost of the high-fidelity function, i.e. 1. The estimated costs of each method with different parameters are shown in Table 2. All gradient-based methods (M1 and M3) showed considerable cost advantages over the surrogate-based M2. These advantages will be further magnified as the problem dimensionality grows. However, the initial points showed strong impacts in the cost and in the final optima for this multi-modal function. As for M2, although the initial surrogate accuracy affected the convergence iteration, the final costs remained similar around 37. Finally, it is clearly noted that the proposed M3 required the lowest cost of 14 and delivered the desired global optimum.

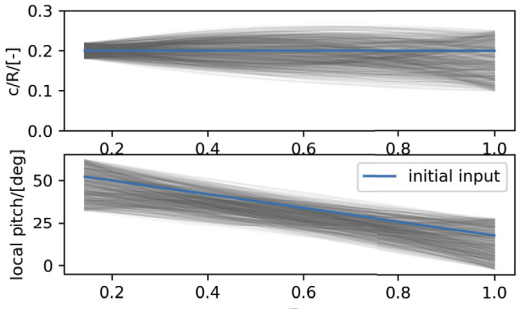

(a) Four-bladed designs.

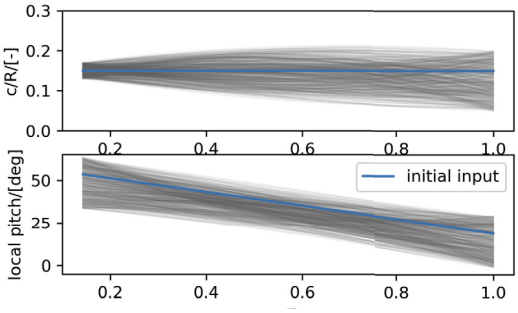

(b) Five-bladed designs.

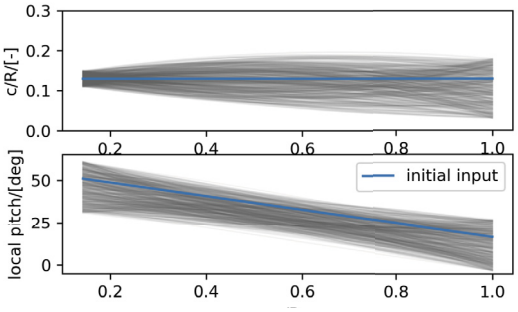

(c) Six-bladed designs.

Fig. 5. Propeller shape parametrisation and sampling with different blade numbers. The grey lines represent sampled shapes. Each family of designs have 1000 sampling points.

4. Multi-fidelity propeller design

4.1. Parametrisation and sampling

For the low-fidelity stage in this study, only blade pitch and chord distributions were considered. This was mostly because the XRotor code [14,15,35] was not able to handle more complex parameters such as blade sectional changes, sweep, anhedral/dihedral angles etc. The later high-fidelity optimisation adopted the same design variables for consistency and for reduced computational cost. The chord distribution was parametrised as $C(r_0) = \sum_i^N B_i c_i r_0$, while the pitch distribution was represented by $\beta(r_0) = \sum_i^N B_i \beta_i r_0$, where $r_0 = r/R$ is the local radial location, B_i is the Bernstein basis function, and β_i and c_i are coefficients used as the design variables. The blade designs adopted the NACA 6412 airfoil and its aerodynamic characteristics were estimated using XFOIL.

Three simple propeller designs, with 4, 5, 6 blades, respectively, were first generated using the simple Blade Element Momentum Theory (BEMT) as the initial input to the low-fidelity stage. These initial designs had the same radius (3.25 m), RPM (396), and solidity (0.078) from the given vehicle specifications. They were all assigned simple rectangular planforms with an estimated linear twist of $-40°$, and the specific chord and pitch were found through the BEMT to meet a target thrust of 39.5 kN (Table 3). Initial attempts were also made to derive ideal designs with minimum induced power through the XRotor inverse design functionality, but the results were highly impractical subject to the objective/constraints (root chord longer than the radius and pitch angle greater than 90 degrees) and hence were abandoned. Note that the solidity values were kept the same here to maintain similar levels of blade loadings, but this constraint was not imposed in the later sampling and optimisation for non-linear shapes.

The three initial geometric inputs were approximated by Bernstein polynomials as shown in Figs. 5(a) to 5(c), using 3 c_i coefficients for the chord and 3 β_i coefficients for the twist distribution, respectively (in total 6 design variables for each shape). Upon these simple initial inputs, the Latin Hypercube Sampling approach was used to generate a total of 3000 potential designs to be evaluated. This is illustrated in Figs. 5(a) to 5(c), with each of the grey transparent lines denoting a new pitch or chord distribution deviated from the initial input. The upper and lower design variable bounds are shown in Table 4. Note that the chord length near the blade root was restricted to avoid root overlapping for high blade numbers.

Table 3
Initial geometry input and design space definitions for the low-fidelity stage.

Planform	rectangular	
Linear twist	-40	deg
Solidity	0.078	
Radius	3.25	m
Blade number	4, 5, 6	
Design variables	chord, twist distributions	
Target thrust	39.5	kN
RPM	396	

Table 4
Upper and lower bounds of the design variable (DV) variations relative to the initial shape.

DV Index	1	2	3
$\Delta\beta_i$/[-]	[-3.18, 1.59]	[-3.18, 1.59]	[-3.18, 1.59]
Δc_i/[-]	[-0.02, 0.02]	[-0.1, 0.1]	[-0.1, 0.05]

4.2. Low-fidelity performance evaluation and filtering

The aerodynamic performance of these potential designs were quickly evaluated using the lower-order code XRotor [14] in an automated manner with the help of bash scripts. All designs assumed the same RPM as per the operational constraints. The computation of all 3000 samples took about 12 minutes on a single 3.2 GHz Intel Xeon CPU. The evaluated performance results of the 4-, 5-,

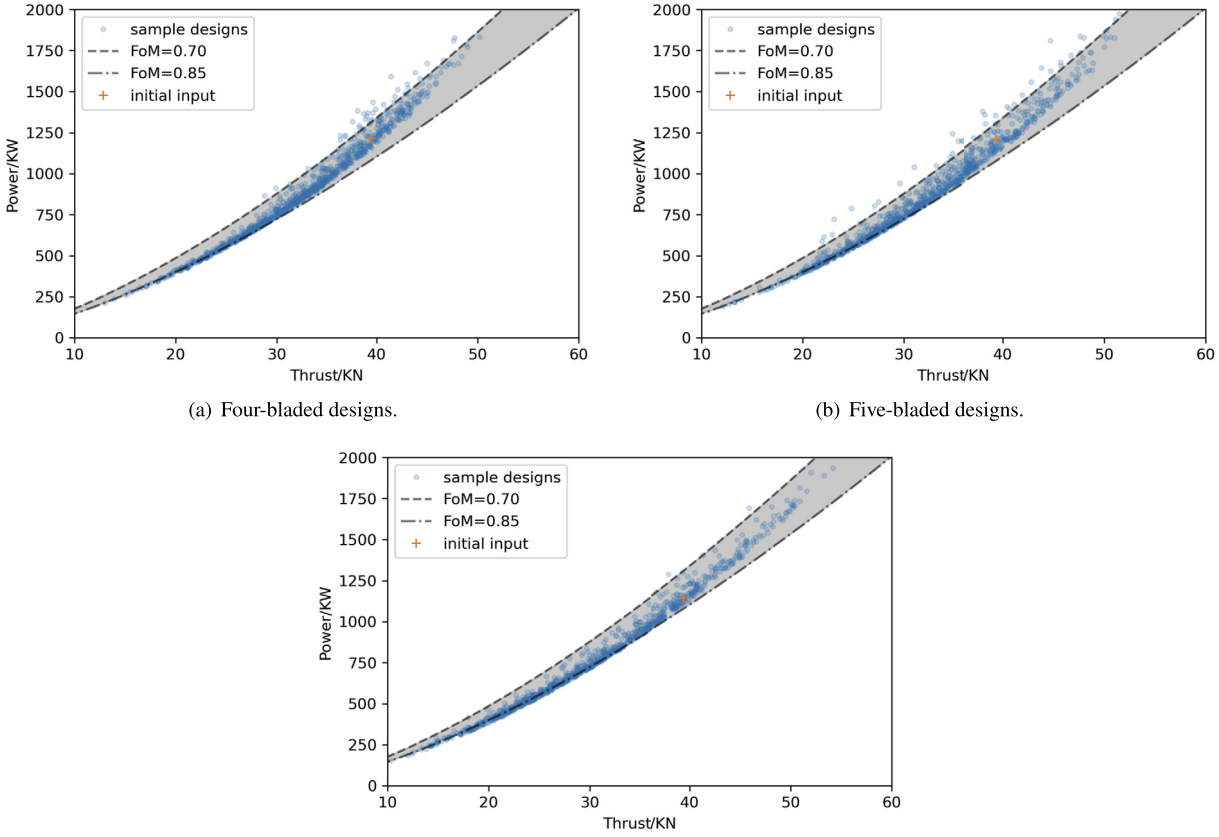

(a) Four-bladed designs.

(b) Five-bladed designs.

(c) Six-bladed designs.

Fig. 6. Low-fidelity performance evaluations of potential designs.

and 6-bladed designs are plotted in Figs. 6(a), 6(b), and 6(c), respectively. Also plotted are the Figure of Merit (FoM) intervals to indicate the aerodynamic efficiencies of designs. It can be observed that, for each family of designs with a different blade number, the distribution of points forms a performance front. As the blade number was increased, the efficiency front of the propeller design was also increased.

The next step was to filter out the propeller designs that resides on the performance front and delivers the target thrust. This was achieved by trimming all potential designs to the target thrust through collective pitch angle changes, then ranking the designs according to their power requirements. This is illustrated in Figs. 7(a) to 7(c) for the three families of designs. The vertical lines represent the thrust target. As can be noted, although all potential designs could be trimmed to provide the target thrust, some required significantly higher power input than others. The red crosses denote the best designs among the potential designs.

Figs. 8(a) to 8(c) present comparisons between the initial designs and selected best designs, in terms of chord and twist distributions. Compared to the simple initial designs, the improved designs showed increased chord around the blade middle span, and reduced chord near the blade tip. Local pitch angles were reduced by about 10 degrees near the blade root and were very slightly increased near the blade tip for all improved designs. Overall, these changes in the blade shapes were expected to shift more aerodynamic loading towards the middle span of the blades.

Table 5 presents more quantitative power comparisons between the initial and best designs while meeting the target thrust requested. It can be noted that the propeller performance improved following the increasing blade numbers for both initial and best designs. The maximum power reduction was about 6.6% relative to the initial input for the five-bladed design. Still, among the improved designs, the six-bladed design showed the best performance with a 3% power reduction relative to the initial input.

Table 5
Performance comparisons between the initial inputs and the low-fidelity optimised designs.

Blade number	Power input/[kW]		Power reduction/[%]
	Initial input	Low-fidelity optimised	
4	1215.3	1152.4	5.2%
5	1207.5	1127.5	6.6%
6	1145.1	1110.9	3.0%

At this stage, the low-fidelity analysis suggested that the six-bladed propeller design is the best choice for the vehicle. However, it should be stressed that these were results from low-fidelity design and analysis tools with higher levels of uncertainty. This is also why optimisation algorithms were not invoked in this stage, as the optimisation may suffer heavily from the uncertainty levels. High-fidelity numerical analyses and optimisation are presented in the next subsection to verify and confirm the performance predictions and to further improve the propeller design.

4.3. High-fidelity analysis and optimisation

The selected optimal designs of Figs. 8(a) to 8(c) from the low-fidelity stage contained the blade shape definitions, but these needed to be converted into 3D shapes in CAD systems and high-quality grids were also necessary for high-fidelity CFD analyses. This was achieved through an automatic mesh generation framework [16,17] of the HMB3 toolkit. Using the blade shape defini-

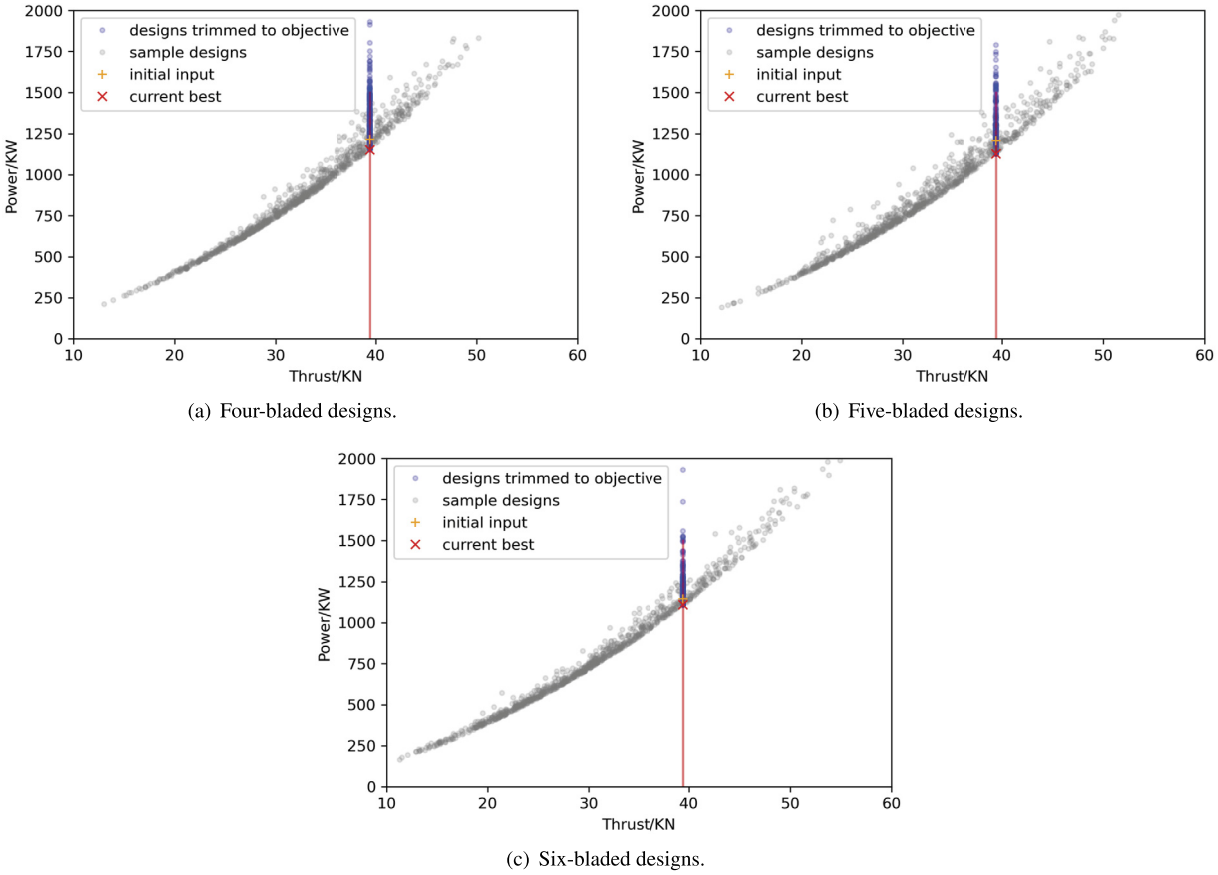

(a) Four-bladed designs.

(b) Five-bladed designs.

(c) Six-bladed designs.

Fig. 7. Selecting the optimum design with the lowest power input in the low-fidelity stage.

tions as the input parameters, geometries and chimera grids were automatically generated through this framework with minor human interventions.

The constructed propeller shapes are shown in Figs. 9(a) to 9(c). The blade sections adopted the same NACA6412 airfoil shape, for simplicity. Fully-structured, body-fitted grids were generated enclosing the blade shapes. The high-fidelity HMB3 simulations were performed in the non-inertial rotating reference frame with the help of periodic and chimera boundaries. For simulations containing only one blade, about 6.3 million cells were used, which was more than sufficient according to previous simulations of propellers in axial flight [20]. The k-ω SST model was adopted for the simulations. The grids were carefully refined based on previous experience [16,20,27], and the y+ values near wall boundaries were kept around 1.

Flow solutions and wake structures of the initial designs at similar thrust levels are shown in Figs. 9(a) to 9(c). The propeller wake structures were clean near the tips for all designs, but the flow solutions became complex near the blade roots. This was due to the low rotational speed, high pitch, high chord, and low sectional thickness at the root. As the blade number was increased, the root solutions improved. It is also interesting to observe that the helical tip vortex stayed closer to the rotor as the blade number increased. The tip vortex descent rates were about $1.67\times10^{-3}R$, $1.39\times10^{-3}R$, and $1.25\times10^{-3}R$ per azimuth degree, respectively, for the four-, five-, and six-bladed designs. This was caused by the lower downwash. It also indicates that a propeller design with many blades may be prone to blade-vortex-interaction (BVI) at low thrust conditions.

Figs. 10(a) to 10(c) present the high-fidelity performance evaluation of the propeller designs in Figs. 8(a) to 8(c) using HMB3 at various RPM settings, along with comparisons against XRotor results correspondingly at the same pitch and RPM settings. The solid lines represent HMB3 results, while the dashed lines represent XRotor results. It is expected that the XRotor results deviate from the HMB3 predictions given their lower-order nature. The XRotor code slightly over-predicted the thrust, torque, and efficiency values for all propeller designs. However, both methods predicted similar performance trends with RPM changes, and the differences between the methods were mainly offsets. It was also confirmed again by the high-fidelity methods that the six-bladed propeller design has the best performance compared to the rest of the tested designs. Later performance analyses and design optimisation hence focused on the six-bladed design.

After high-fidelity performance evaluation of the design output from the low-fidelity stage, the next step is to further improve the design through high-fidelity methods. In the current work, gradient-based shape optimisation methods were adopted with the help of adjoint computations. The low-fidelity design output was used as the initial design input to the optimisation loop.

As stated earlier, the optimisation study focused on the six-bladed design. The objective was to reduce the propeller torque, i.e. the power consumption, while maintaining the thrust. The design variables considered were again the blade pitch and chord distributions along the blade span. Note that the HMB3 adjoint is capable of arbitrary shape changes [20,21], but the current work focused only on the twist and chord for consistency and for reduced computational costs. The HMB3 adjoint treated the twist and chord changes as arbitrary surface deformations regardless, through the IDW mesh deformation [19,33,34] module. This was a classic non-linear optimisation problem subject to equality constraints. The optimisation problem was hence formulated as follows:

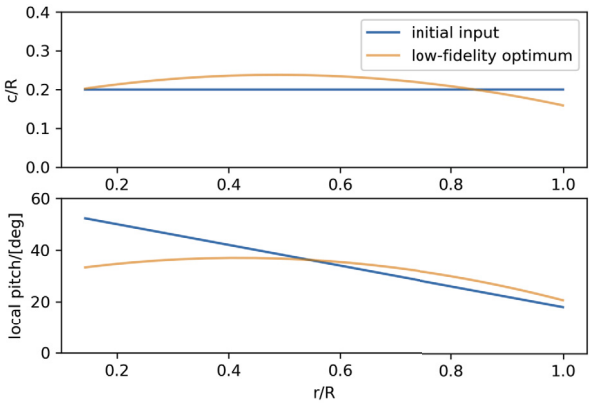

(a) Four-bladed designs.

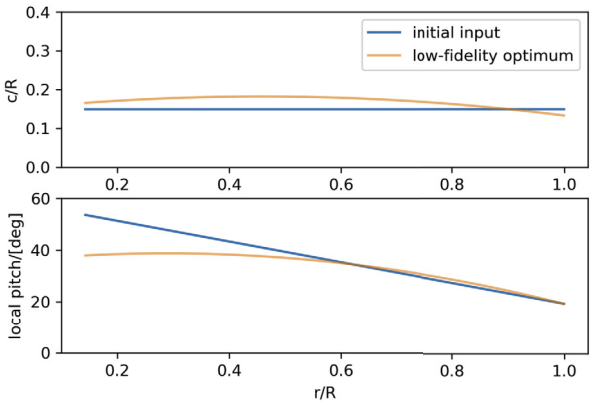

(b) Five-bladed designs.

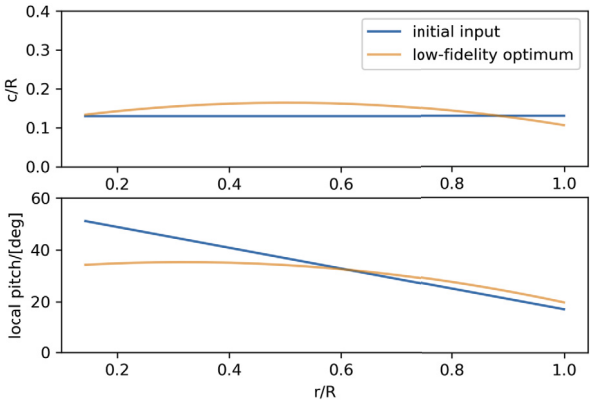

(c) Six-bladed designs.

Fig. 8. Comparisons between the initial rectangular inputs and the selected low-fidelity optimum designs among the potential designs.

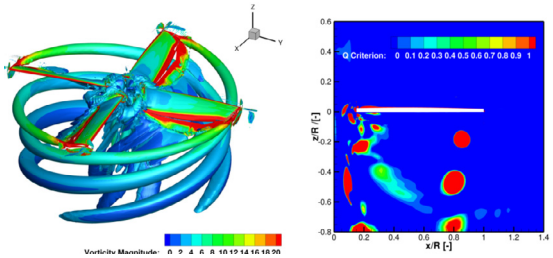

(a) Four-bladed designs.

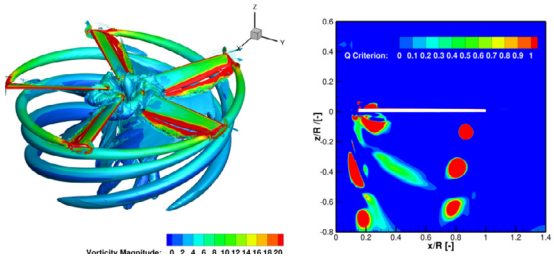

(b) Five-bladed designs.

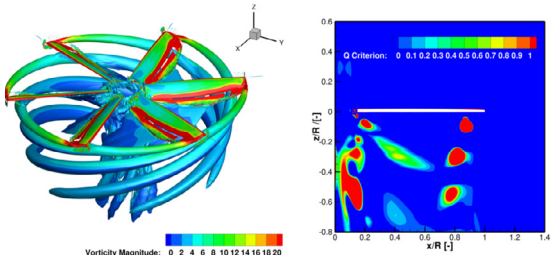

(c) Six-bladed designs.

Fig. 9. Flow solutions and wake structures of the propeller designs, illustrated using q-criterion iso-surfaces coloured by vorticity magnitude.

Find: $min(Q(\beta_i, c_i))$

by varying: (β_i, c_i)

subject to: $|T(\beta_i, c_i) - T_0| \leq \epsilon$

where $Q(\beta_i, c_i)$ and $T(\beta_i, c_i)$ are the propeller torque and thrust, respectively. T_0 is a reference thrust level that should be maintained, i.e. the target thrust. (β_i, c_i) are again Bernstein polynomial coefficients representing, respectively, the twist and chord distributions. In the present work, 8 design variables (4 for the chord and 4 for the twist) were used. Upon the initial input

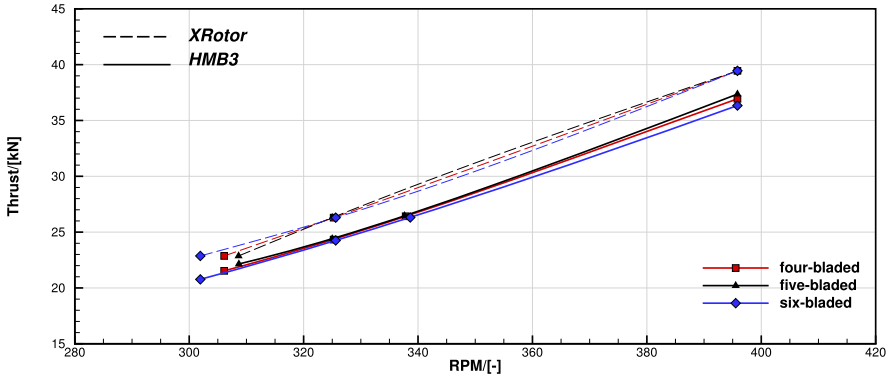

(a) Thrust.

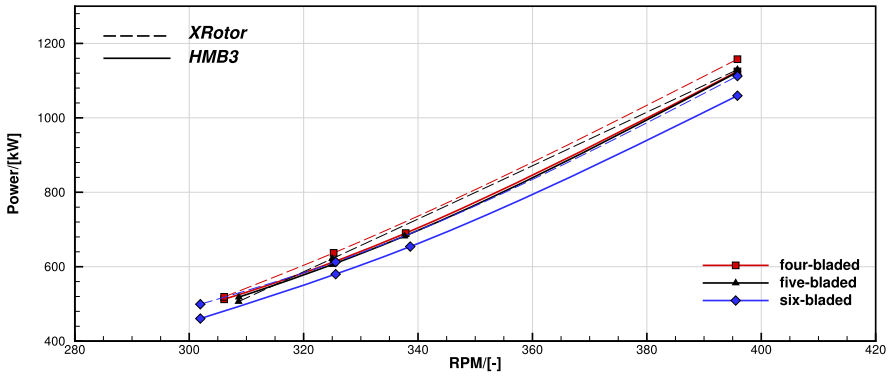

(b) Power.

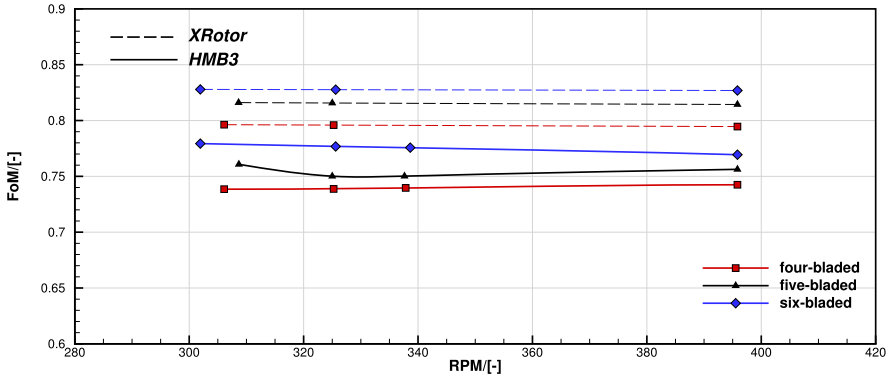

(c) Efficiency.

Fig. 10. Comparisons between low-fidelity and high-fidelity aerodynamic performance evaluations of the propeller designs.

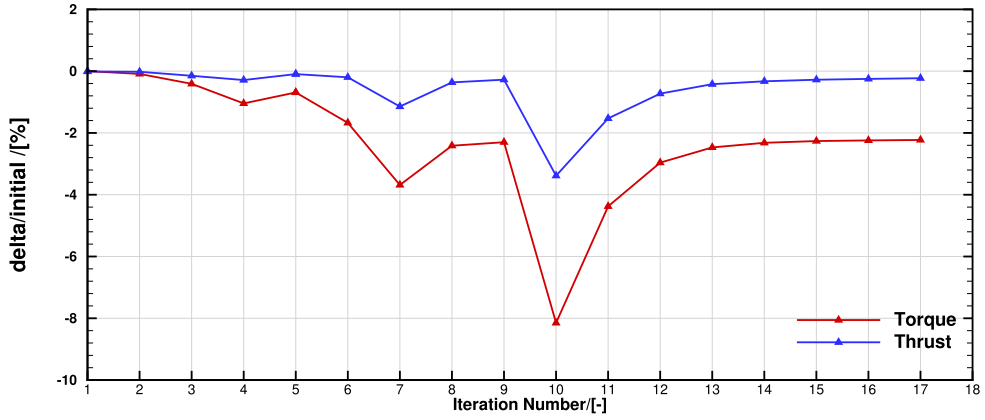

Fig. 11. High-fidelity blade shape optimisation convergence history.

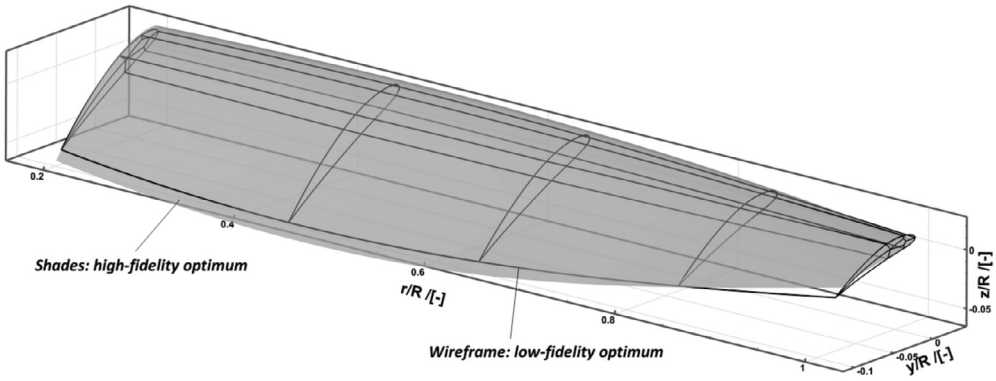

(a) Comparisons between the high-fidelity optimised and the low-fidelity optimum blade shapes. The green shape is the initial low-fidelity design, while the red is the high-fidelity optimised design.

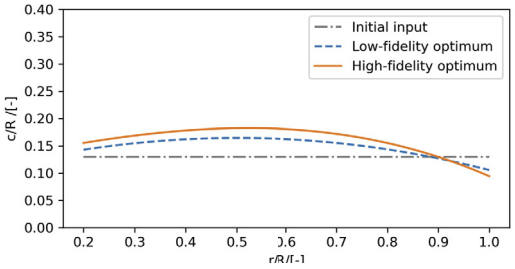

(b) Chord distribution changes.

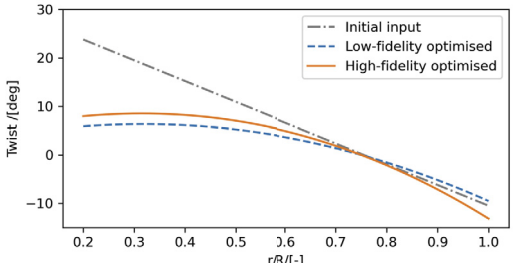

(c) Twist distribution changes.

Fig. 12. Changes in the blade twist and chord distributions brought by the high-fidelity optimisation.

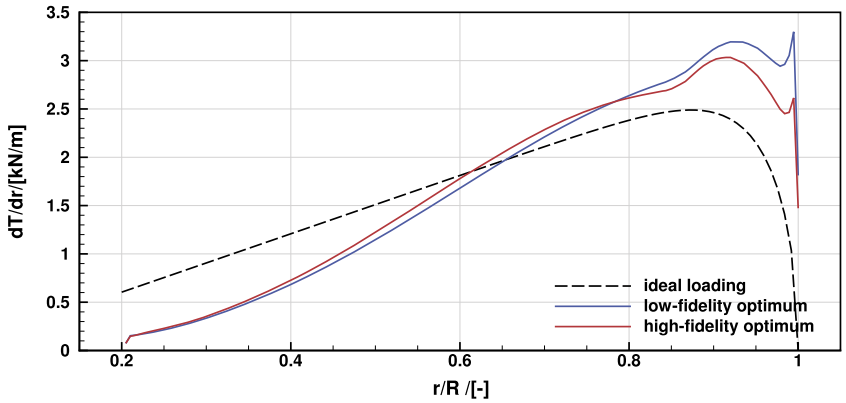

Fig. 13. Blade loading comparisons between the low-fidelity/high-fidelity optima (both computed through HMB3) and the ideal loading with Prandtl's tip loss correction [42] at the same pitch, RPM, and thrust.

(β_{i0}, c_{i0}), the chord design variables were bounded to vary within a $\Delta c_i = \pm 0.05$, while the twist design variables were bounded within $\Delta\beta_i = \pm 0.5$. The number of design variables were also increased to allow for more delicate designs using high-fidelity methods. Note that physical scopes of the design variables and objectives/constraints remain unchanged.

The optimisation was performed using a medium grid of 4 million cells per blade, but the optimisation results were later verified using refined grids, and differences in the performance predictions were minor. The optimisation convergence history is shown in Fig. 11. The optimisation converged within 17 iterations. Each design iteration required approximately 18 minutes using 64 Intel Xeon CPU cores (3.2 GHz). The final results represent about 2.2% decrease in the torque and negligible changes in the thrust (about 0.2% decrease). The torque reduction corresponded to about 15 kW power reduction for the propeller, while changes in the thrust were only about 50 N.

Comparisons between the optimised and initial shapes are shown in Figs. 12(a), 12(b), and 12(c). The blade chord was increased from the root up until $r/R = 0.9$. Towards the tip, the blade chord was reduced. The optimisation also increased the overall blade twist. The local pitch angle near the root was increased, while the tip pitch angle was reduced. These geometric changes led to reduced tip loading and increased loading towards the midspan and the root as shown in Fig. 13.

Comparisons of thrust loading along the blade radius between the low-fidelity optimum, high-fidelity optimum at the same pitch angle, RPM, and thrust are presented in Fig. 13, which were both extracted from the HMB3 computations. Also included is the ideal blade loading (with Prandtl's tip loss correction) [42] correspond-

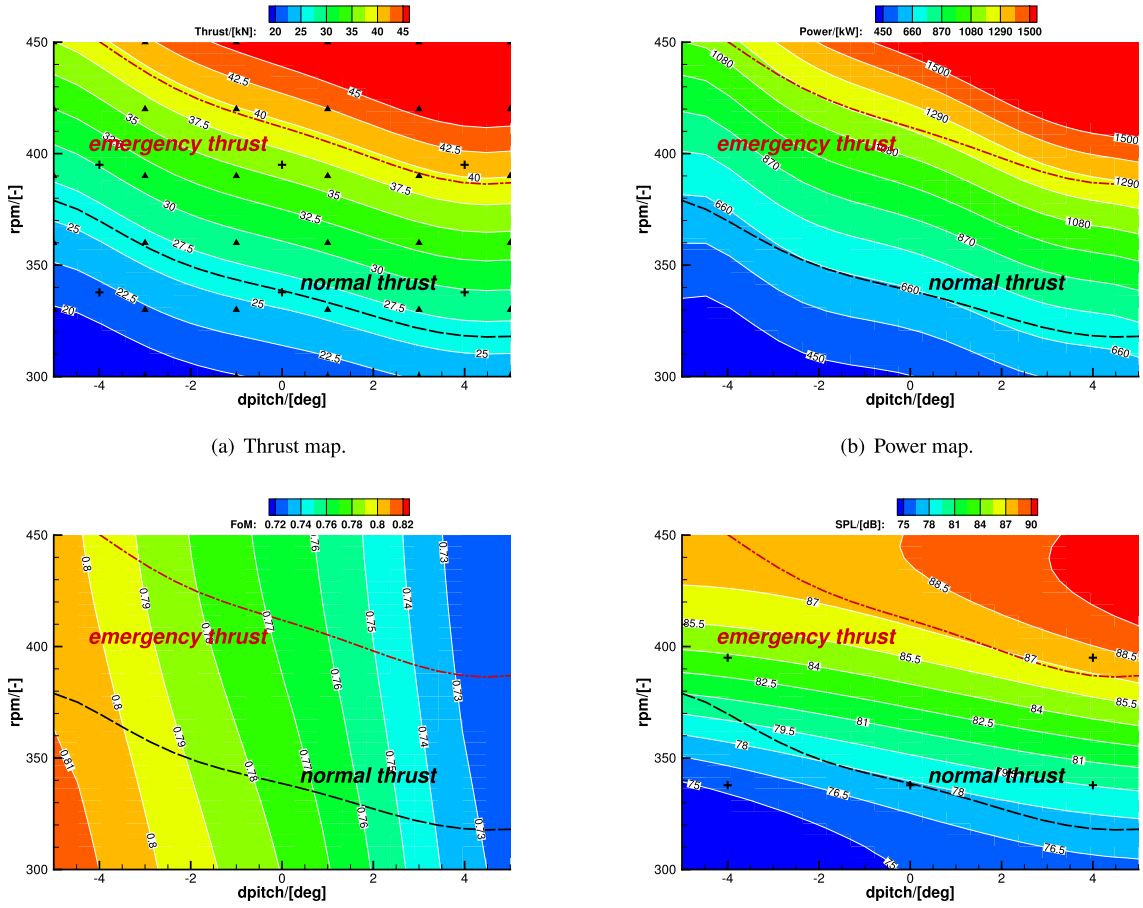

(a) Thrust map.

(b) Power map.

(c) Efficiency map.

(d) Near-field acoustics (Kriging on high-fidelity data only.)

Fig. 14. Performance map of the optimised propeller design subject to pitch and RPM regulations, represented by multi-fidelity Kriging approximations (+: high-fidelity samples, Δ: low-fidelity samples).

ing to the minimum induced power with a uniform downwash. Compared to the ideal loading, the low-fidelity optimum had higher loading towards the blade tip. The high-fidelity shape optimisation made the blade loading closer to the ideal distribution by shifting more loading from the tip towards the mid-span and root sections. This was consistent with the chord and twist changes. However, changes near the root sections were minimum. The root loading could be improved by altering the sectional shape, which was not done in the present optimisations. Further performance evaluations at various operating conditions of the optimised propeller design are presented in later sections.

5. Propeller pitch-RPM performance maps

For conventional pitch-regulated propellers, the performance output (either thrust, torque, or acoustic/elastic responses) at a given flight condition, can be written as a function $f(\beta)$ subject to the collective pitch setting β. Future eVTOL vehicles, however, can have both pitch and RPM regulations for their propellers thanks to the adoption of electrical motors. The introduction of variable RPM control expands the function to $f(\beta, RPM)$ which certainly unlocks a much larger performance space and adds further operation flexibility, yet the problem dimensionality has also been increased. This poses challenges for the propeller performance evaluation, as well as, the determination of operating conditions. This section hence describes the pitch-RPM performance map approach that has effectively facilitated the development of the Skybus vehicle [43] in terms of performance evaluation and operating condition determination.

5.1. Construction and verification

The performance maps presented in this work are essentially response surfaces constructed upon the low-fidelity and high-fidelity propeller performance datasets. This work examined three different response surface constructions, i.e. Kriging models upon low-fidelity XRotor data, Kriging models upon high-fidelity HMB3 data, and multi-fidelity Kriging model integrating both datasets. In Fig. 14(a), the "+" represent the 6 high-fidelity sampling points while the "Δ" signs denote the 36 orthogonal low-fidelity sampling points. Note that the acoustic response surface (Fig. 14(d)) was constructed using high-fidelity data only, as the low-fidelity tool was unable to provide accurate acoustic predictions.

To quantify the response surface error, Table 6 presents the Leave-One-Out (LOO) verifications of the three models in terms of standard deviations of relative errors at each sampling point. Note that the LOO verification for the multi-fidelity model was carried on the high-fidelity points. The efficiency error was not included here as it was derived upon the thrust and torque predictions, and may be affected by error cancellation. Overall, all models in Table 6 showed small errors near or below 1% for thrust. The 1.78% acoustic error from high-fidelity Kriging was also seen as reasonable. The error for power prediction was slightly larger near 3%, which was expected as the torque values are typically more sensitive to

Table 6
Leave-One-Out (LOO) verification of the response surface construction denoted by the standard deviation (σ) of relative errors (δ) at each point.

Response Surfaces	$\sigma(\delta T)/[\%]$	$\sigma(\delta P)/[\%]$	$\sigma(\delta SPL)/[\%]$
Low-fidelity Kriging (36)	1.19	3.73	-
High-fidelity Kriging (6)	1.32	3.24	1.78
Multi-fidelity Kriging (36+6)	0.97	2.6	-

Table 7
Verification of response surface predictions at an extra high-fidelity HMB3 sampling point at $(0°, 325.6)$ in terms of relative errors.

Response Surfaces	$\delta T/[\%]$	$\delta P/[\%]$	$\delta FoM/[\%]$	
Low-fidelity Kriging (36)	9.07	6.64	6.49	-
High-fidelity Kriging (6)	0.2	1.2	-1.3	1.3
Multi-fidelity Kriging (36+6)	-0.02	-0.22	0.05	-

the modelling errors. Among all three models, the multi-fidelity model clearly showed the lowest construction error.

To quantify the model prediction accuracy, Table 7 further presents the relative prediction errors at an extra high-fidelity HMB3 sampling point at $(0°, 325.6)$. The low-fidelity response surface showed large prediction errors for the thrust (9%), power (6.64%), and efficiency (6.49%). The high-fidelity Kriging, although with only 6 sampling points, effectively reduced the predictions near 1%. The multi-fidelity model clearly offered the most accurate performance predictions at point, with all errors well below 0.5%. Further analysis for the Skybus design hence focused on this multi-fidelity model as shown in Figs. 14(a) to 14(c).

5.2. Performance/operating condition interpolation

Figs. 14(a) to 14(d) present performance maps of the optimised propeller design output of the high-fidelity stage, in terms of thrust, power, efficiency (Figure of Merit), and Sound Pressure Levels (SPL). Note that the acoustic data (SPL) was included here to highlight the multi-disciplinary feature of the performance map. Acoustic data was extracted directly from the HMB3 simulations at a microphone point 10 m away, 2 m below the propeller centre of rotation. In these figures, the horizontal axes represent blade pitch changes, while the vertical axes are propeller RPM values. The contour maps represent the specific performance data within the pitch/RPM intervals.

These maps present clear illustrations of various performance variations subject to pitch/RPM changes. The propeller aerodynamic and acoustic performance at a specific pitch/RPM combination can be easily interpolated from the performance maps. For the range studied, the propeller thrust and power increase with the RPM and pitch increases as shown in Figs. 14(a) and 14(b). The propeller efficiency, however, is mostly dominated by the pitch changes as shown in Fig. 14(c). Lower pitch angles lead to higher efficiency while the RPM has less effects. As for the acoustics in Fig. 14(d), the RPM shows stronger impacts as the noise increases quickly with increased RPM, but higher pitch also leads to stronger noise. More performance outputs, e.g. blade loadings or aero-elastic responses, can also be extracted and added to the performance maps to provide more comprehensive performance evaluation.

During the development of the Skybus vehicle [43], these maps have been used to effectively determine the operating conditions through graphic optimisation. The dashed lines in Figs. 14(a) correspond to the target thrust, and every pitch/RPM combination on this line would fulfil the thrust requirement. Overlapping this target thrust line with the other performance maps, as shown in Figs. 14(b) to 14(d), the potential power, efficiency, and acoustic outputs subject to this thrust constraint are highlighted. It then becomes an intuitive graphic-based process to select the best operating condition subject to the equality thrust constraint. Depending on the specific objectives, the operating condition can be easily determined. For instance, if one wishes to have maximum efficiency, the propeller collective pitch angle should be reduced while the RPM should be maintained around 380 according to Fig. 14(c). On the contrary, if acoustic reduction is of the primary concern, such as in the Skybus case, the propeller pitch angle should be increased while the RPM should be reduced according to the acoustic performance map in Fig. 14(d). This approach has successfully supported the high-fidelity simulation of the entire Skybus vehicle as reported in Ref [43].

6. Conclusions

This work presented propeller design analysis process to support the development of the heavy-lift eVTOL Skybus. A multi-fidelity approach was first proposed and demonstrated for the propeller design of the Skybus vehicle. To further analyse the propeller performance subject to pitch and RPM controls, pitch/RPM performance maps were proposed and composed using multi-fidelity Kriging surrogate models.

The proposed multi-fidelity propeller design framework combines fast low-fidelity scans of the design space, and high-fidelity gradient-based optimisations starting from the low-fidelity optima, thereby delivering improved final designs subject to constraints. The methodology was first evaluated using the benchmark Styblinski-Tang function. Compared to classic gradient-based or multi-fidelity surrogate-based approaches, the proposed multi-fidelity method delivered the global optimum at the lowest cost.

The proposed framework was then applied to the eVTOL propeller design. Simple rectangular propeller designs with four, five, and six blades were first evaluated, and the six-bladed designs were found to be the most efficient. For the six-bladed propeller design of the Skybus vehicle, the low-fidelity stage brought a performance improvement of 3%, while the high-fidelity optimisation further boosted the performance by about 2.2% with negligible penalties in thrust. The optimised blade design had a non-linear chord distribution with increased chord length in the mid-span region. A non-linear twist of about 20 degrees towards the blade tip was also derived. These design modifications shifted the blade loading towards the blade mid-span and brought the overall loading distribution closer to the ideal loading, improving the overall performance.

This work also reported the pitch-RPM performance map tool that has efficiently supported the Skybus vehicle development. For both pitch and RPM regulated eVTOL propellers, the performance map visualises and correlates multiple scopes, and can support quick performance estimation. The multi-fidelity construction and verification of the performance maps were presented in detail. The multi-fidelity Kriging model offered accurate predictions with errors less than 0.3% for the validation point. The graphic-based selection of operating conditions for the Skybus vehicle through the performance map were also demonstrated.

Future work will continue to explore multi-point and multi-objective design tasks using the proposed design framework. Future design and analysis of the Skybus will be continued as well.

Declaration of competing interest

There is no conflict of interests.

Data availability

Data will be made available on request.

Acknowledgement

The authors would like to acknowledge GKN Aerospace for their support and approval of publication. The project was also supported by the UK Engineering and Physical Sciences Research Council. The support received from the University of Glasgow, EPSRC Impact Acceleration Account (EP/R511705/1) is gratefully acknowledged

References

[1] A. Filippone, G. Barakos, Rotorcraft systems for urban air mobility: a reality check, Aeronaut. J. (2020) 1–19.
[2] J. Graham, Definition of a rotorcraft, https://www.aerosociety.com/Assets/Docs/Publications/SpecialistPapers/Definition_of_a_Rotorcraft.pdf, February 2013. Special paper of the Royal Aeronautical Society, available on-line.
[3] R.A. Ormiston, Revitalising advanced rotorcraft research - and the compound helicopter, Aeronaut. J. 120 (2016) 83–129, https://doi.org/10.1017/aer.2015.5.
[4] D.E. Myers, Matrix formulation of co-kriging, J. Int. Assoc. Math. Geol. 14 (3) (1982) 249–257.
[5] M.C. Kennedy, A. O'Hagan, Predicting the output from a complex computer code when fast approximations are available, Biometrika 87 (1) (2000) 1–13.
[6] S. Gaggero, G. Vernengo, D. Villa, A two-fidelity level approach for marine propeller design, in: The 9th Conference on Computational Methods in Marine Engineering (Marine 2021), 2022.
[7] N. Mourousias, A. Malim, B.G. Marinus, M. Runacres, Assessment of multi-fidelity surrogate models for high-altitude propeller optimization, in: AIAA AVIATION 2022 Forum, 2022, p. 3752.
[8] P.Z. Korondi, Multi-fidelity surrogate-assisted design optimisation under uncertainty for computationally expensive aerospace applications, Ph.D. thesis, 2021.
[9] A. Molina-Cristóbal, P.R. Palmer, B.A. Skinner, G.T. Parks, Multi-fidelity simulation modelling in optimization of a submarine propulsion system, in: 2010 IEEE Vehicle Power and Propulsion Conference, IEEE, 2010, pp. 1–6.
[10] J. Martins, A. Ning, Engineering Design Optimization, 2021.
[11] D.J. Mavriplis, Discrete adjoint-based approach for optimization problems on three-dimensional unstructured meshes, AIAA J. 45 (4) (2007) 741–750, https://doi.org/10.2514/1.22743.
[12] M. Biava, M. Woodgate, G.N. Barakos, Fully implicit discrete-adjoint methods for rotorcraft applications, AIAA J. 54 (2) (2015) 735–749, https://doi.org/10.2514/1.J054006.
[13] T. Zhang, G. Barakos, High-fidelity numerical investigation of ducted propeller aero-dynamics/acoustics and adjoint-based design optimisation, in: 77th Annual Forum of the Vertical Flight Society, 2021.
[14] M. Drela, H. Youngren, XROTOR user guide, http://web.mit.edu/drela/Public/web/xrotor/, 2003. Available on-line.
[15] H. Youngren, M. Drela, S. Sanders, Ducted fan design code, http://web.mit.edu/drela/Public/web/dfdc/, December 2005. Ducted Fan Design Code homepage, available on-line.
[16] T. Zhang, G.N. Barakos, High-fidelity CFD validation and assessment of ducted propellers for aircraft propulsion, J. Am. Helicopter Soc. 66 (1) (2021) 1–28, https://doi.org/10.4050/JAHS.66.012008.
[17] T. Zhang, G.N. Barakos, Development of simulation tools for high fidelity analysis of compound rotorcraft, in: AIAA Scitech 2020 Forum, 2020, p. 1258.
[18] R.J. Higgins, A. Zarev, G.N. Barakos, R.B. Green, Numerical investigation of a two-bladed propeller inflow at yaw, J. Aircr. 57 (2) (2020) 292–304.
[19] M. Biava, G.N. Barakos, Optimisation of ducted propellers for hybrid air vehicles using high-fidelity CFD, Aeronaut. J. 120 (1232) (2016) 1632–1657, https://doi.org/10.1017/aer.2016.78.
[20] T. Zhang, G.N. Barakos, High-fidelity numerical analysis and optimisation of ducted propeller aerodynamics and acoustics, Aerosp. Sci. Technol. 113 (2021) 106708, https://doi.org/10.1016/j.ast.2021.106708.
[21] T. Zhang, G.N. Barakos, Toward vehicle-level optimization of compound rotorcraft aerodynamics, AIAA J. 60 (3) (2022) 1937–1957.
[22] R. Steijl, G.N. Barakos, K. Badcock, A framework for CFD analysis of helicopter rotors in hover and forward flight, Int. J. Numer. Methods Fluids 51 (8) (2006) 819–847, https://doi.org/10.1002/fld.1086.
[23] A. Antoniadis, D. Drikakis, B. Zhong, G. Barakos, R. Steijl, M. Biava, L. Vigevano, A. Brocklehurst, O. Boelens, M. Dietz, et al., Assessment of CFD methods against experimental flow measurements for helicopter flows, Aerosp. Sci. Technol. 19 (1) (2012) 86–100, https://doi.org/10.1016/j.ast.2011.09.003.
[24] R. Steijl, G. Barakos, CFD analysis of complete helicopter configurations–lessons learnt from the GOAHEAD project, Aerosp. Sci. Technol. 19 (1) (2012) 58–71, https://doi.org/10.1016/j.ast.2011.01.007.
[25] D. Han, V. Pastrikakis, G.N. Barakos, Helicopter flight performance improvement by dynamic blade twist, Aerosp. Sci. Technol. 58 (2016) 445–452.
[26] A.J. Garcia, G.N. Barakos, Numerical simulations on the ERICA tiltrotor, Aerosp. Sci. Technol. 64 (2017) 171–191, https://doi.org/10.1016/j.ast.2017.01.023.
[27] T. Zhang, G. Qiao, D. Smith, G. Barakos, A. Kusyumov, Parametric study of aerodynamic performance of equivalent ducted/un-ducted rotors, Aerosp. Sci. Technol. 117 (2021) 106984, https://doi.org/10.1016/j.ast.2021.106984.
[28] G. Chirico, G.N. Barakos, N. Bown, Propeller installation effects on turboprop aircraft acoustics, J. Sound Vib. 424 (2018) 238–262.
[29] M. Jarkowski, M. Woodgate, G. Barakos, J. Rokicki, Towards consistent hybrid overset mesh methods for rotorcraft CFD, Int. J. Numer. Methods Fluids 74 (8) (2014) 543–576, https://doi.org/10.1002/fld.3861.
[30] F. Menter, Two-equation eddy-viscosity turbulence models for engineering applications, AIAA J. 32 (8) (1993) 1598–1605, https://doi.org/10.2514/3.12149.
[31] D. Kraft, Algorithm 733: TOMP-Fortran modules for optimal control calculations, ACM Trans. Math. Softw. 20 (3) (1994) 262–281, https://doi.org/10.1145/192115.192124.
[32] S. Johnson, The NLopt nonlinear-optimization package, https://nlopt.readthedocs.io/en/latest/, retrieved 3 Sep 2020.
[33] D. Shepard, A two-dimensional interpolation function for irregularly-spaced data, in: Proceedings of the 1968 23rd ACM National Conference, 1968, pp. 517–524.
[34] R.J. Higgins, A. Jimenez-Garcia, G.N. Barakos, N. Bown, High-fidelity computational fluid dynamics methods for the simulation of propeller stall flutter, AIAA J. 57 (12) (2019) 5281–5292.
[35] M. Drela, H. Youngre, Axisymmetric analysis and design of ducted rotors, http://web.mit.edu/drela/Public/web/dfdc/DFDCtheory12-31.pdf, December 2005. Ducted Fan Design Code theory, available on-line.
[36] M. Drela, XFOIL: an analysis and design system for low Reynolds number airfoils, in: Low Reynolds Number Aerodynamics, Springer, 1989, pp. 1–12.
[37] J. Sacks, S.B. Schiller, W.J. Welch, Designs for computer experiments, Technometrics 31 (1) (1989) 41–47, https://doi.org/10.1080/00401706.1989.10488474.
[38] L. Le Gratiet, Multi-fidelity Gaussian process regression for computer experiments, Ph.D. thesis, Université Paris-Diderot-Paris VII, 2013.
[39] M.A. Bouhlel, J.T. Hwang, N. Bartoli, R. Lafage, J. Morlier, J.R.R.A. Martins, A Python surrogate modeling framework with derivatives, Adv. Eng. Softw. (2019) 102662, https://doi.org/10.1016/j.advengsoft.2019.03.005.
[40] M. Styblinski, T.-S. Tang, Experiments in nonconvex optimization: stochastic approximation with function smoothing and simulated annealing, Neural Netw. 3 (4) (1990) 467–483.
[41] D.R. Jones, M. Schonlau, W.J. Welch, Efficient global optimization of expensive black-box functions, J. Glob. Optim. 13 (4) (1998) 455–492, https://doi.org/10.1023/A:1008306431147.
[42] W. Johnson, Helicopter Theory, Courier Corporation, 2012.
[43] T. Zhang, G. Barakos, M. Foster, et al., Multi-fidelity aerodynamic and acoustic design and analysis of a heavy-lift eVTOL, in: 48th European Rotorcraft Forum, Winterthur, Switzerland, 2022.